

Minimum Weight Euclidean t -spanner is NP-Hard *

Paz Carmi and Lilach Chaitman-Yerushalmi

Department of Computer Science,
Ben-Gurion University of the Negev, Israel

Abstract. Given a set P of points in the plane, an Euclidean t -spanner for P is a geometric graph that preserves the Euclidean distances between every pair of points in P up to a constant factor t . The weight of a geometric graph refers to the total length of its edges. In this paper we show that the problem of deciding whether there exists an Euclidean t -spanner, for a given set of points in the plane, of weight at most w is NP-hard for every real constant $t > 1$, both whether planarity of the t -spanner is required or not.

1 Introduction

Consider a weighted graph $G = (V, E)$ with a weight function $w : E \rightarrow \mathbb{R}^+$ over the edges. For any two vertices $u, v \in V$, we denote the weight of the shortest path between u and v in G by $\delta_G(u, v)$. Given a spanning subgraph G' of G , we define the *dilation* of G' (with respect to G) to be the value

$$\max_{u, v \in V} \frac{\delta_{G'}(u, v)}{\delta_G(u, v)}.$$

Given a real value $t > 1$, a t -spanner of G is a spanning subgraph G' with dilation at most t with respect to G . Thus, the the shortest-path distances in G' approximate shortest-path distances in the underlying graph G within an approximation ratio t . Typically, G is a dense graph with $\Omega(n^2)$ edges and the t -spanner G' is desired to be sparse, preferably having only a linear number of edges.

Spanners have been studied in many different settings. The various settings differ from one another in the characterization of the underlying graph G , such as different topologies and specific weight functions over the edges, in the value of the required dilation t , and in the properties of the spanner G' , such as planarity. We concentrate on the setting where the underlying graph is geometric. In our context, a graph $G = (P, E)$ is called geometric graph or Euclidean graph if its vertex set P is a set of points in the plane and every edge $\{p, q\}$ in E is the line segment \overline{pq} , weighted by the Euclidean distance between its endpoints $|pq|$.

* Research is partially supported by the Lynn and William Frankel Center for Computer Science and by grant 680/11 from the Israel Science Foundation (ISF).

Moreover, the underlying graphs we consider are complete graphs and therefore we refer to their t -spanners as spanners of the point set P . There is a vast body of literature on t -spanners in this geometric setting (see [10] for a comprehensive survey of the area).

The weight of a geometric graph is defined as the sum of the lengths of its edges. The weight is a good measure of the cost of building the network; thus, it is often desirable to have spanners with low weight. Since any spanner must connect all the points, the weight of a t -spanner is bounded from below by the weight of a minimum spanning tree $MST(P)$. Chandra et al. [4] have presented a greedy algorithm for constructing a t -spanner in $O(n^3 \log n)$ time, which has been proved by Das et al. [5, 7] to have a weight of size $O(\text{wt}(MST(P)))$. The constant factor depends on the value t . A more efficient algorithm that computes the greedy spanner in $O(n^2 \log n)$ was later developed by Bose et al. [1]. A fast implementation of a variant of the greedy algorithm that maintains the $O(\text{wt}(MST(P)))$ weight and runs in $O(n \log^2 n)$ time has been developed by Das and Narasimhan [6]. Those spanners approximate the weight of the minimum t -spanners within a constant factor that depends on t , but they are not necessarily optimal. We address the following decision problem and appropriate optimization problem.

Problem 1. *The Low WEIGHT t -SPANNER (LWSt) decision problem:*

Input: *A set P of points in the plane and a constant $w > 0$.*

Output: *Whether there exists an Euclidean t -spanner for P of weight at most w .*

Problem 2. *The MINIMUM WEIGHT t -SPANNER (MWSt) problem:*

Input: *A set P of points in the plane.*

Output: *A minimum weight Euclidean t -spanner for P .*

We show that for every real value $t > 1$ the MWSt and the LWt problems are NP-hard. This is done by a reduction from the PARTITION problem, defined as follows.

Problem 3. *The PARTITION problem:*

Input: *A set $X = \{x_1, \dots, x_n\}$ of n positive integers with even $\sum_{x \in X} x = R$.*

Output: *Whether there exists a subset $X' \subset X$ such that $\sum_{x \in X'} x = R/2$.*

Klein and Kutz [9] have proved that the DILATION GRAPH and the PLANE DILATION GRAPH problems are NP-hard by a reduction from the PARTITION problem as well. The DILATION GRAPH problem (resp. the PLANE DILATION GRAPH problem) ask whether there exists a t -spanner (resp. a plane t -spanner) with at most m edges for a given set of points P , an integer m , and a real value $t > 1$. Note that here t is part of the input. Their reduction returns an instance with $t = 7$ for every instance of the PARTITION problem.

Various minimization problems of different parameters such as weight and number of edges of t -spanners have been proved to be NP-hard. Cai et al. have proved in [2] that for $t \geq 4$ the problem of determining the existence of a t -spanner with at most m edges for an unweighted graph is NP-hard. This implies that the problem of finding a minimum t -spanner of weight at most w for a

weighted graph is NP-hard as well. However, no similar conclusions regarding the NP-hardness of the problem in our geometric setting, i.e., of MWSt and LWSt can be deduced.

Brandes and Handke have introduced a related problem of finding a minimum weight planar t -spanner for weighted graphs. They have established it is NP-hard for $t > 1$ by modifying the proof in [3] of the NP-hardness of the tree t -spanner problem for weighted graphs. A variation of the minimum weight planar t -spanner problem adjusted to the geometric setting discussed in this paper should restrict the underlying graph to be the complete Euclidean graph. The appropriate decision and optimization problems are defined as follows.

Problem 4. *The Low WEIGHT PLANE t -SPANNER (LWPSt) decision problem:*

Input: A set P of points in the plane and a constant $w > 0$.

Output: Whether there exists an Euclidean plane t -spanner for P of weight at most w .

Problem 5. *The MINIMUM WEIGHT PLANE t -SPANNER (MWPSSt) decision problem:*

Input: A set P of points in the plane and a constant $w > 0$.

Output: A minimum weight Euclidean Plane t -spanner for P .

The reductions presented in this paper prove that the LWPSt and MWPSSt problems are NP-hard for every $t > 1$.

Regarding geometric graphs, Gudmundsson and Smid consider in [8] the problem of deciding whether a given geometric graph (not necessarily the complete graph) contains a t -spanner with at most m edges and prove it is NP-hard for every $t > 1$. This implies that the problem of finding a spanning graph of a given geometric graph with at most m edges and minimum dilation is also NP-hard. In this paper we consider the minimality of spanners in terms of weight and not in terms of the number of edges as addressed in this problem. In addition, we restrict the underlying graph to be the complete Euclidean graph. Therefore, we suggest the following variation.

Problem 6. *The MINIMUM DILATION GRAPH (MDG) decision problem:*

Input: A set P of points in the plane and a constant $w > 0$.

Output: A minimum dilation Euclidean graph for P of weight at most w .

From the NP-hardness of the LWSt problem proved in this paper, we may deduce the NP-hardness of the MDG problem.

Note that although all the problems presented here refer to the restricted case where the underlying graph is the complete Euclidean graph, our results, obviously, apply to the modified problems where the underlying graph is a general geometric graph (not necessarily the complete graph) and where the underlying graph is a general weighted graph (not necessarily geometric).

The rest of the paper is organized as follows. In section 2 we define new terms and make some technical observations to be used in the reductions proofs. The reductions are presented and proved in section 3. First the reductions idea is outlined and later described in more detail, where the reduction for $t \geq 2$ is given in subsection 3.1 and the reduction for $1 < t < 2$ is given in subsection 3.2.

2 Definitions and Technical Lemmas

Definition 1. Given a path $\{p, s, q\}$, a t -shortcut refers to the addition of the edge $\{p, q\}$ and possibly a removal of one of the edges $\{p, s\}$ or $\{s, q\}$, as long as the obtained graph is a t -spanner for $\{p, s, q\}$.

Definition 2. Given a path $Q = \{p, s, q\}$ and a t -shortcut that results in a graph G ,

- the benefit of the t -shortcut is defined as $\delta_Q(p, q) - \delta_G(p, q) = |ps| + |sq| - |pq|$,
- the cost of the t -shortcut is defined as $\text{weight}(G) - \text{weight}(Q)$, and
- the efficiency of the t -shortcut is defined as the ratio between its benefit and its cost, i.e., $\frac{\delta_Q(p, q) - \delta_G(p, q)}{\text{weight}(G) - \text{weight}(Q)}$.

Lemma 1. Given two paths $Q = \{p, s, q\}$ and $Q' = \{p', s', q'\}$, such that $|ps| = |sq|$, $|p's'| = |s'q'|$, and $\angle(psq) < \angle(p's'q')$ (see Fig. 1), and let e and e' denote the two efficiencies of the most efficient t -shortcuts in Q and Q' , respectively, then $e > e'$.

Proof. By the fact that $\angle(psq) < \angle(p's'q')$, we conclude

$$\frac{|ps|}{|pq|} > \frac{|p's'|}{|p'q'|}. \quad (1)$$

The most efficient t -shortcut in Q obviously includes the addition of $\{p, q\}$, and if $|ps| + |pq| \leq t|qs|$ it also includes the removal of $\{q, s\}$ (or alternatively $\{p, s\}$).

If indeed $|ps| + |pq| \leq t|qs|$, then

$$\begin{aligned} e &= \frac{|ps| + |sq| - |pq|}{|pq| - |sq|} = 1/\left(\frac{|pq|}{|ps|} - 1\right) - 1 \\ &>^{(1)} 1/\left(\frac{|p'q'|}{|p's'|} - 1\right) - 1 = \frac{|p's'| + |s'q'| - |p'q'|}{|p'q'| - |s'q'|} \geq e'. \end{aligned}$$

Thus, $e > e'$ as required.

Otherwise, $|ps| + |pq| > t|qs|$ and therefore

$$1 + \frac{|pq|}{|qs|} > t \stackrel{(1)}{\Rightarrow} 1 + \frac{|p'q'|}{|q's'|} > t \Rightarrow |p's'| + |p'q'| > t|q's'|$$

and neither $\{q, s\}$ nor $\{q', s'\}$ are removed from the respective most efficient t -shortcuts. Hence,

$$\begin{aligned} e &= \frac{|ps| + |sq| - |pq|}{|pq|} = \frac{2|ps|}{|pq|} - 1 \\ &>^{(1)} \frac{2|p's'|}{|p'q'|} - 1 = \frac{|p's'| + |s'q'| - |p'q'|}{|p'q'|} = e', \end{aligned}$$

and $e > e'$ as required. □

Lemma 2. Given a path $Q = \{p, s, q\}$ such that $|ps| < |sq|$, $\angle(psq) > \frac{\pi}{2}$, and $-\cos(\angle(psq)) \geq \frac{1}{k+1}$ for a positive value k (see Fig. 1), the efficiency of a t -shortcut in Q is less than k .

Proof. Let e denote the efficiency of the most efficient t -shortcut in Q and let $a = |sq|$, $b = |ps|$, $c = |pq|$, and $\alpha = \angle(psq)$. Note that $e \leq \frac{a+b-c}{c-a}$.

Since it is given that $-\cos \alpha \geq \frac{1}{k+1}$, by the cosines law we have

$$c^2 = a^2 + b^2 - 2ab \cos \alpha \geq a^2 + b^2 + \frac{2ab}{k+1} > a^2 + \frac{b^2}{(k+1)^2} + \frac{2ab}{k+1} = \left(a + \frac{b}{k+1}\right)^2.$$

Thus, we receive $c > a + \frac{b}{k+1}$ and hence $e \leq \frac{a+b-c}{c-a} < k$.

□

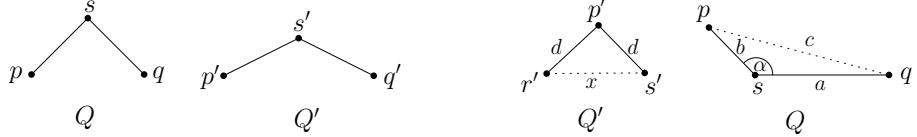


Fig. 1. On the left, the paths Q and Q' as defined in Lemma 1; on the right, the paths Q and Q' as defined in Lemma 2 and Corollary 1.

Corollary 1. Given four points $\{r', s', s, q\}$ on a line and two additional points p and p' , such that $|r'p'| = |p's'|$, $|ps| < |sq|$, and the two edges (p, s) and (p', s') are parallel, we consider the two paths $Q' = \{r', p', s', s\}$ and $Q = \{p, s, q\}$ (see Fig. 1). Let e' and e denote the two efficiencies of the most efficient t -shortcuts in Q' and Q , respectively, then $e' > e$.

Proof. Let $x = |r's'|$ and $d = |r'p'| = |p's'|$, then $e' \geq \frac{2d-x}{x} = 2d/x - 1$. By the sines law, $\angle(rps) = 2 \arcsin(\frac{x}{2d})$. Since the angle $\angle(psq)$, denoted by α , equals to the exterior angle of the isosceles triangle $\triangle(rps)$, we have $\alpha = 2 \arcsin(\frac{x}{2d}) + \pi/2 - \arcsin(\frac{x}{2d})/2 = \pi - \arccos(\frac{x}{2d})$ and $\cos(\alpha) = -\frac{x}{2d}$. By Lemma 2, we receive $e < 2d/x - 1 \leq e'$. □

3 The reduction

In this section we show that the LWSt decision problem (Problem 1) and the MWSt optimization problem (Problem 2) are NP-hard for every constant $t > 1$. We prove the NP-hardness of the LWSt problem by a reduction from the PARTITION problem (Problem 3). The NP-hardness of the MWSt problem follows from an obvious reduction from the appropriate decision problem (the LWSt problem). We propose different reductions for $1 < t < 2$ and for $t \geq 2$; however, both follow the same core ideas.

Given an instance $X = \{x_1, \dots, x_n\}$ with $\sum_{x \in X} x = R$ for the PARTITION problem, both reductions output a weight w and a set P of linear size in n that is composed of points distributed along a path connecting n triples of isosceles triangles' vertices. Each isosceles triangle gadget is associated with a value x_i among x_1, \dots, x_n and the length of its edges are derived from the value x_i . The distances between adjacent points on the path connecting all triangles are derived from the sum R .

In order to prove the correctness of the reduction, namely, that a subset $X' \subset X$ with $\sum_{x \in X'} x = R/2$ exists iff there exists a t -spanner for P of weight at most w , we use the same core method in both cases. First, we observe that the minimum connected graph over P forms a path $\{p = p_1, \dots, p_m = q\}$ of weight (length) $w - c \cdot R = t|pq| + l \cdot R$. Meaning, a total shortening of $l \cdot R$ should be performed on the path at the cost of at most $c \cdot R$.

We show that any shortening can be considered as a set of independent t -shortcuts and that the most *efficient* (as defined in Definition 2) t -shortcuts are those that involve adding a base of an isosceles triangle gadget and their efficiency is exactly l/c . We refer to such t -shortcuts as *gadget t -shortcuts*. Therefore, no other shortcut may be applied and the distance between p and q can be decreased to $t|pq|$ iff there exists a set of gadget t -shortcuts with a total benefit of $l \cdot R$. Our construction ensures that a set of gadget t -shortcuts with a total benefit of $l \cdot R$ exists iff there exists a subset $X' \subset X$ such that $\sum_{x \in X'} x = R/2$ and applying those t -shortcuts creates a t -spanner for P . Moreover, the t -spanner obtained by applying those t -shortcuts is a plane graph. Therefore, our reductions also prove the NP-hardness of the LWPSt problem (Problem 4) and the MWPSSt problem (Problem 5).

The reduction for $t \geq 2$ is a bit simpler and therefore described first, in the next subsection, followed by the reduction for $1 < t < 2$ in subsection 3.1.

3.1 The reduction for $t \geq 2$

Given a valid input for the PARTITION problem $X = \{x_1, \dots, x_n\}$ with $\sum_{x \in X} x = R$, the reduction outputs a valid input for the SWt problem (P, w) as elaborated next. The points of P are placed roughly along 3 sides of an axis-parallel rectangle, the left side, the right side, and the top edge as illustrated in Fig. 2. The left and right sides are of length $R/2(t(n+2) - n - 2\frac{1}{3})$; each is sampled by $\lceil t(n+2) - n - 2\frac{1}{3} \rceil$ points with regular spacing of $R/2$ (possibly except for the bottom point). The top edge is actually a horizontal component of width $R(n+2)$, consists of n isosceles triangle gadgets and connecting points. The top edge itself is sampled by $2n$ endpoints of the n isosceles triangles' bases, located among $n+1$ segments of length R , halved by a middle point. The additional vertex of each triangle is located above the rectangle. The i -th isosceles triangle gadget has sides of length $\frac{5}{6}x_i$, a base of length x_i , and subtended angle of $2\arcsin(\frac{3}{5}) < \pi/2$. Overall, we get $|P| = 4n + 3 + 2\lceil t(n+2) - n - 2\frac{1}{3} \rceil$. One can verify that P is of polynomial size in n and can be represented in polynomial size in n and $\log R$, and thus can be generated in polynomial time.

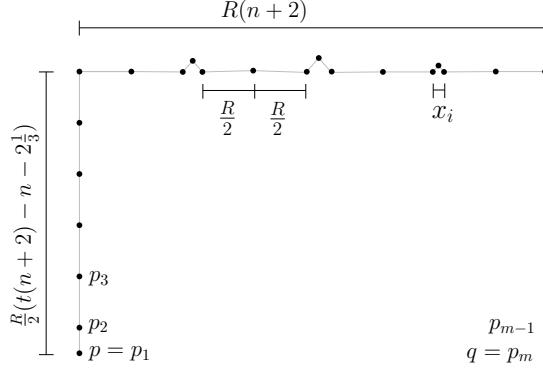


Fig. 2. The set of points P and its minimum connected graph as defined in the reduction for $t \geq 2$.

The weight bound w is defined a bit differently for $2 \leq t < 2\frac{1}{5}$ and $t \geq 2\frac{1}{5}$. For $2 \leq t < 2\frac{1}{5}$ we define $w = R(t(n+2) + \frac{5}{6})$ and for $t \geq 2\frac{1}{5}$, $w = R(t(n+2) + \frac{5}{12})$. We are now ready to prove the correctness of the reduction; namely, that a subset $X' \subset X$ with $\sum_{x \in X'} x = R/2$ exists iff there exists a t -spanner for P of weight at most w .

First, observe that the minimum connected graph over P forms a path $\{p = p_1, \dots, p_m = q\}$ as depicted in Fig. 2 of weight $2 \cdot R/2(t(n+2) - n - 2\frac{1}{3}) + R(n+1) + \frac{5}{3}R = R(t(n+2) + \frac{1}{3})$. Since $|pq| = R(n+2)$, this path is $R/3$ units longer than a legal t -spanning path from p to q ; however, it is lighter then the given weight w . For $2 \leq t < 2\frac{1}{5}$ the remaining weight is $R/2$ and for $t \geq 2\frac{1}{5}$ the remaining weight is $R/12$. Therefore, a total shortening of $R/3$ should be performed on the path at the cost of at most $R/2$ or $R/12$ for $2 \leq t < 2\frac{1}{5}$ and $t \geq 2\frac{1}{5}$ respectively.

Assuming there exists a t -spanner for P of weight at most w , let E^+ and E^- be the sets of edges over P that should be added and removed from $MST(P)$, respectively, in order to obtain a t -spanner $G = (P, E)$ that minimizes $weight(G) \leq w$. We make some observations regarding the edge set E^+ .

Claim 1. *The edge set E^+ contains only edges of length at most R .*

Proof. Since the initial graph is a minimum spanning tree and G must be connected, $|E^-| \leq |E^+|$; moreover, there exists an injective function from E^- to E^+ that maps every edge $e^- \in E^-$ to a longer edge $e^+ \in E^+$ whose addition creates a cycle in $MST(P)$ and enables the removal of e^- . Assume towards contradiction that E^+ contains an edge e^+ longer than R . Since all edges in $MST(P)$ are of length at most $\max\{R/2, \max_{1 \leq i \leq n} \{5x_i/6\}\} \leq (*) R/2$, we have that

$$\sum_{e \in E^+} weight(e) - \sum_{e \in E^-} weight(e) > R/2$$

and thus $weight(G) > w$. Inequality $(*)$ holds under the legitimate assumption that all the elements in X are smaller than $R/2$. \square

According to the above claim, only edges of length at most R can be added to $MST(P)$. There are exactly three types of such edges (while ignoring edges that overlaps subpaths in $MST(P)$, which are useless for shortening the distance between p and q):

1. A base edge of an isosceles triangle gadget on the top edge.
2. An edge connecting a top vertex of an isosceles triangle on the top side with a point on the top side itself that closes a cycle of size three.
3. An edge connecting the left (resp. right) edge with the top edge and closes a cycle of size three.

Observation 1. *Any two edges in E^+ close edge-disjoint cycles in $MST(P)$.*

Proof. The only pair of edges that may violate this condition is a pair of an edge of type 1 and an edge of type 2. Assume towards contradiction that there are two such edges, e_1 of type 1 and e_2 of type 2, in E^+ . Due to the minimality of $weight(G)$ there is necessarily a t -spanning path connecting two points in G that contains e_2 . Note that between every two points on the top side of the rectangle there is a t -spanning path in $MST(P)$; hence, e_2 is necessarily contained in a t -spanning path connecting two points on different sides of the rectangle that passes through the top side. According to the same arguments, the same holds for e_1 . Meaning, e_1 and e_2 are both contained in the shortest path connecting the endpoints of the top side of the rectangle (see Fig. 3). However, replacing e_1 and e_2 with the induced subgraph of $MST(P)$ over the endpoints of e_1 and e_2 results in a shorter path. Therefore, the addition of e_1 and e_2 to E^+ is not necessary for creating any t -spanning path and thus contradicts the minimality of $weight(G)$. \square

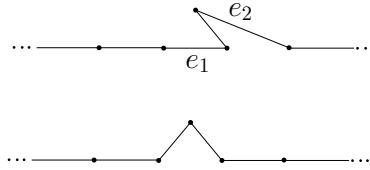


Fig. 3. On the top, part of the t -spanning path between two points in G containing e_1 and e_2 , as discussed in the proof of Observation 1. On the bottom, part of the t -spanning path between the same two points in $MST(P)$.

By Observation 1 we conclude that G can be obtained from $MST(P)$ by applying a set of t -shortcuts in edge-disjoint subpaths of $MST(P)$. As we have already observed, each t -shortcut involves an addition of an edge of type 1, 2, or 3 and possibly a removal of an edge in the closed cycle. The addition of an edge of type 1, together with a removal of an isosceles triangle side for $t \geq 2\frac{1}{5}$, we refer to as a *gadget t*-shortcut.

Claim 2. *The most efficient possible t -shortcuts are gadget t -shortcuts. Their efficiency is $\frac{2}{3}$ for $2 \leq t < 2\frac{1}{5}$ and 4 for $t \geq 2\frac{1}{5}$.*

Proof. First, note that the removal of a triangle side in a gadget t -shortcut for $t \geq 2\frac{1}{5}$ is indeed possible in the aspect of dilation. Now let us examine the efficiency of a gadget t -shortcut. The benefit of the i -th gadget t -shortcut is $\frac{5}{3}x_i - x_i = 2x_i/3$. For $2 \leq t < 2\frac{1}{5}$ its cost is x_i and for $t \geq 2\frac{1}{5}$ its cost is $1\frac{5}{6}x_i - \frac{5}{3}x_i = x_i/6$; hence, its efficiency is indeed $2/3$ and 4 for $2 \leq t < 2\frac{1}{5}$ and $t \geq 2\frac{1}{5}$, respectively.

Next, we observe every conceivable t -shortcut and show it is less efficient than a gadget t -shortcut.

- Addition of an edge of type 2 forms an obtuse triangle. By Corollary 1, this t -shortcut is less efficient than the gadget t -shortcut.
- Addition of an edge of type 3 forms an isosceles right triangle and by Lemma 1 the gadget t -shortcut is more efficient.

□

According to Claim 2, only the gadget t -shortcuts have efficiency equal to the ratio between the required shortening of the path between p and q and the remaining weight. Therefore, those are the only shortcuts that can be applied. Moreover, in order to achieve a $R/3$ shortening of the path connecting p and q without exceeding the weight bound, a set of gadget t -shortcuts with a total benefit of exactly $R/3$ should be applied.

Lemma 3. *A subset $X' \subset X$ with $\sum_{x \in X'} x = R/2$ exists iff there exists a t -spanner for P of weight at most w .*

Proof. [\Leftarrow] Assume towards contradiction that there exists a t -spanner for P of weight at most w , however, a subset $X' \subset X$ with $\sum_{x \in X'} x = R/2$ does not exist. According to our construction, this implies that a set of gadget t -shortcuts with a total benefit of exactly $R/3$ does not exist. By the aforementioned analysis, this means that there is no graph over P of weight at most w in which there exists a path connecting p and q of length at most $t|pq|$. This contradicts the existence of a t -spanner for P of weight at most w .

[\Rightarrow] Assume that a subset $X' \subset X$ with $\sum_{x \in X'} x = R/2$ exists. Let $G = (P, E)$ denote the graph obtained by applying the gadget t -shortcuts in the triangles that correspond to the elements in X' . We show that G admits a t -spanner for P , i.e., for every two points u and v in P , $\delta_G(u, v) \leq t|uv|$. We have already shown that there exists a t -spanning path between p and q in G . Next we consider all other pairs of points $\{u, v\} \neq \{p, q\}$ in P :

1. One of $\{u, v\}$ is on the right side of the rectangle and the other is on the left side: $\delta_G(u, v) \leq \delta_G(p, q)$ and $|uv| \geq |pq|$, hence $\delta_G(u, v)/|uv| \leq \delta_G(p, q)/|pq| = t$.

2. One of $\{u, v\}$ is on the right or left side of the rectangle and the other is on the top side: Assume w.l.o.g. that u is on the left side and v is on the top side (see Fig. 4). Let a denote the left-top corner point of the rectangle. We have

$$\begin{aligned}
\delta_G(u, v)/|uv| &= (\delta_G(u, a) + \delta_G(a, v))/|uv| \\
&\stackrel{(*)}{\leq} \max\{(|ua| + |av|(\frac{4}{5} + \frac{1}{5} \cdot 1\frac{5}{6}/\frac{1}{2}))/|uv|, \\
&\quad (|ua| + |av|(\frac{2}{3} + \frac{1}{3} \cdot \frac{5}{3}))/|uv|\} \\
&\leq (|ua| + 1\frac{8}{15}|av|)/\sqrt{|ua|^2 + |av|^2} \\
&<^{(**)} 2 \leq t|uv|
\end{aligned}$$

Inequality $(*)$ holds under the assumption that all the elements in X are smaller than $R/2$ and due to the notion that every triangle gadget base is followed by two edges of length $R/2$ on the left. Inequality $(**)$ holds since the values of the function $(x + 1\frac{8}{15}y)/\sqrt{x^2 + y^2}$ are smaller than 2 for every x and y .

3. Both u and v are on the top side of the rectangle: As noted before, if u and v are the endpoints of a side of an isosceles triangle gadget we have $\delta_G(u, v) \leq t|uv|$. Let u' and v' be the projections of u and v on the line containing the top side of the rectangle (see Fig. 4). By the same arguments as in the previous case, we have

$$\delta_G(u, v) \leq (\frac{2}{3} + \frac{1}{3} \cdot 1\frac{5}{6}/\frac{1}{2})|u'v'| < 2|uv| \leq t|uv|.$$

4. Both u and v are on the left or the right side of the rectangle: We have $\delta_G(u, v) = |uv| < t|uv|$.

□

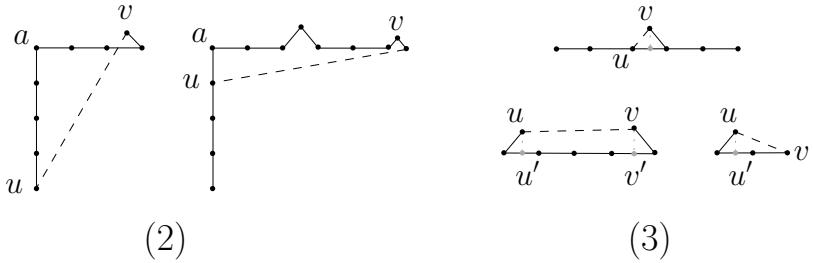


Fig. 4. On the left, illustrations for case 2 in the proof of Lemma 3; on the right, illustrations for case 3 in the same proof.

Overall, we proved a proper reduction for $t \geq 2$.

3.2 The reduction for $1 < t < 2$

We reduce from the PARTITION problem and define the reduction output (P, w) as follows. The weight bound is $w = R(n+t+\frac{3}{2}+(n+\frac{3}{2})\frac{3t}{2})$ and the set of points P resembles the one defined for $t \geq 2$. We still have a horizontal component, and left and right sides; however, those are no longer perpendicular to the horizontal component. The top right angles are enlarged by $\alpha_t = \arcsin(\frac{2}{3t^2} + \frac{1}{3t})$ and the rectangle shape turns into an isosceles trapezoid shape (see Fig. 5).

As before, the horizontal component is composed of a horizontal segment of length $R(n+2)$ sampled by $2n$ endpoints of n isosceles triangles' bases, located among $n+1$ segments of length R halved by a middle point. However, here the i -th isosceles triangle has sides of length $\frac{t}{2}x_i$, a base of length x_i , and thus an angle of $2\arcsin(1/t)$.

The right and left sides are of length $R/2(n+\frac{3}{2})\frac{3t}{2}$; each is sampled by $\lceil(n+\frac{3}{2})\frac{3t}{2}\rceil$ points with regular distances of $R/2$ (possible except for the bottom point).

Overall, we get $|P| = 4n + 3 + 2\lceil(n+\frac{3}{2})\frac{3t}{2}\rceil$. We denote by p and q the leftmost and rightmost points in P , respectively, and by p' and q' the leftmost and rightmost points of the horizontal component, respectively.

The reduction as presented above might not be computable by a Turing machine in polynomial time, since there might be points in P with coordinates whose representation requires large number of bits. This issue is resolved by a minor change as described in the end of this subsection.

We prove the correctness of the reduction, namely, that a subset $X' \subset X$ with $\sum_{x \in X'} x = R/2$ exists iff there exists a t -spanner for P of weight at most w , by following the same guidelines as for the proof for $t \geq 2$.

Note that the minimum connected graph over P forms a path $Q = \{p = p_1, \dots, p_m = q\}$ as depicted in Fig. 5 of length

$$2R/2(n+\frac{3}{2})\frac{3t}{2} + (n+1)R + tR = R(n+t+1+(n+\frac{3}{2})\frac{3t}{2}).$$

Thus, the remaining weight is $R/2$. We would like to examine how much should this path be shortened in order for it to be a proper t -spanning path between p and q . A proper t -spanning path between p and q is of length at most

$$\begin{aligned} t|pq| &= t(2\frac{R}{2}(n+\frac{3}{2})3t/2 \sin \alpha_t + (n+2)R) \\ &= Rt(R(n+\frac{3}{2})\frac{3t}{2} \sin(\arcsin(\frac{2}{3t^2} + \frac{1}{3t})) + n+2) \\ &= Rt((n+\frac{3}{2})\frac{3t}{2}(\frac{2}{3t^2} + \frac{1}{3t} + n+2)) \\ &= R((n+\frac{3}{2})\frac{3t}{2}(1 - (\frac{2}{3} - \frac{2}{3t})) + tn + 2t) \\ &= R((n+\frac{3}{2})\frac{3t}{2} + n + \frac{3}{2} + \frac{t}{2}). \end{aligned}$$

Therefore, a total shortening of $(t-1)R/2$ should be performed on the path at the cost of at most $R/2$.

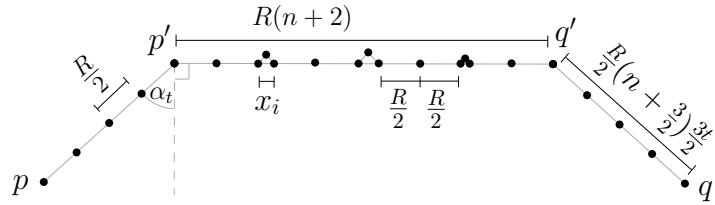


Fig. 5. The set of points P is depicted in black and its minimum connected graph is depicted in gray as defined in the reduction for $1 < t < 2$.

Assuming there exists a t -spanner for P of weight at most w , let E^+ and E^- be the sets of edges over P that should be added and removed from $MST(P)$, respectively, in order to obtain a t -spanner $G = (P, E)$ that minimizes $weight(G) \leq w$. By Claim 1 (which, by the same arguments, holds for this reduction as well), E^+ contains only edges of length at most R . There are exactly four types of such edges (while ignoring edges that overlap subpaths in $MST(P)$, which are obviously redundant):

1. A base edge of an isosceles triangle gadget on the horizontal component.
2. An edge connecting a top vertex of an isosceles triangle on the horizontal component with a point on the horizontal segment and closes a cycle of size three.
3. An edge connecting the left (resp. right) side with the horizontal component and closes a cycle of size three.

One can verify that Observation 1 holds here as well and we have that G can be obtained from $MST(P)$ by applying a set of t -shortcuts that involve an addition of an edge of type 1, 2, or 3 and possibly removal of an edge in the closed cycle. Here, the term *gadget t-shortcut* refers to the addition of an edge of type 1. No removal of an edge is possible for $1 < t < 2$.

Claim 3. *The most efficient possible t -shortcuts are gadget t -shortcuts. Their efficiency is $t - 1$.*

Proof. First we examine the efficiency of a gadget t -shortcut. The benefit of the i -th gadget t -shortcut is $tx_i - x_i = (t-1)x_i$, its cost is x_i , and hence its efficiency is indeed $t-1$.

We now examine every conceivable t -shortcut and show that it is less efficient than a gadget t -shortcut.

- Addition of an edge of type 2 forms an obtuse triangle. By Corollary 1, this t -shortcut is less efficient than the gadget t -shortcut.
- Addition of an edge of type 3 forms an isosceles triangle with angle $\pi/2 + \arcsin(\frac{2}{3t^2} + \frac{1}{3t})$. This angle is greater than $2\arcsin(1/t)$ for every $t > 1$ and by Lemma 1, we conclude that the gadget t -shortcut is more efficient.

□

By Claim 3, in order to achieve a $(t-1)R/2$ shortening of the path connecting p and q , a set of gadget t -shortcuts with a total benefit of exactly $(t-1)R/2$ should be applied. This notion, together with the following Lemma, serve us in the proof of Lemma 5.

Lemma 4. *Let $\triangle(psq)$ be an isosceles triangle with $|ps| = |sq|$ and let q' be a point on \overline{sq} such that $|pq| \geq |pq'|$, then $\frac{|ps|+|sq|}{|pq|} \geq \frac{|ps|+|sq'|}{|pq'|}$.*

Proof. Let $\gamma = \angle(psq)$ and $\beta = \angle(pq'q)$, then by the sines law, $|pq| = |ps|2\sin(\frac{\gamma}{2})$, $|sq'| = |ps|\sin(\beta - \gamma)/\sin(\beta)$, and

$$|pq'| = |pq|\sin(\frac{\pi}{2} - \frac{\gamma}{2})/\sin(\beta) = |ps|2\sin(\frac{\gamma}{2})\cos(\frac{\gamma}{2})/\sin(\beta).$$

Thus,

$$\begin{aligned} \frac{|ps|+|sq|}{|pq|} &\geq \frac{|ps|+|sq'|}{|pq'|} \\ \Leftrightarrow \frac{2|ps|}{|ps|2\sin(\frac{\gamma}{2})} &\geq \frac{|ps|+|ps|\sin(\beta - \gamma)/\sin(\beta)}{|ps|2\sin(\frac{\gamma}{2})\cos(\frac{\gamma}{2})/\sin(\beta)} \\ \Leftrightarrow \frac{\cos(\frac{\gamma}{2})}{\sin(\beta)} &\geq 1 + \frac{\sin(\beta - \gamma)}{\sin(\beta)} \\ \Leftrightarrow \cos(\frac{\gamma}{2}) - \sin(\beta) &\geq \sin(\beta - \gamma) \geq 0. \end{aligned}$$

The last inequality indeed holds for every $0 < \gamma, \beta < \pi$. \square

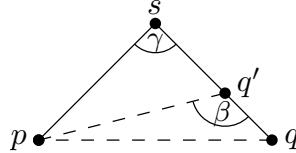


Fig. 6. Illustration of Lemma 4.

Lemma 5. *A subset $X' \subset X$ with $\sum_{x \in X'} x = R/2$ exists iff there exists a t -spanner for P of weight at most w .*

Proof. $[\Leftarrow]$ Assume towards contradiction that there exists a t -spanner for P of weight at most w , however, a subset $X' \subset X$ with $\sum_{x \in X'} x = R/2$ does not exist. According to our construction, this implies that a set of gadget t -shortcuts with a total benefit of exactly $(t-1)R/2$ does not exist. This means there is no graph over P of weight at most w that contains a path connecting p and q of length at most $t|pq|$ in contradiction to the existence of a t -spanner for P of weight at most w .

[\Rightarrow] Assume that a subset $X' \subset X$ with $\sum_{x \in X'} x = R/2$ exists. Let $G = (P, E)$ denote the graph obtained by applying the gadget t -shortcuts in the triangles that correspond to the elements in X' . We show that G admits a t -spanner for P , i.e., for every two points u and v in P , $\delta_G(u, v) \leq t|uv|$. As we have already observed, there exists a t -spanning path between p and q in G . Next we consider all other pairs of points $\{u, v\} \neq \{p, q\}$ in P :

1. One of $\{u, v\}$ is on the right side of the trapezoid and the other is on the left side (see Fig. 7): Assume w.l.o.g. that u is on the left side and v is on the right side. Let s be the intersection point of the extensions of the segments pp' and qq' (see Fig. 7). Let $\epsilon = |p's| + |sq'| - \delta_G(p', q')$, we show that $(|us| + |sv|)/|uv| \leq (|ps| + |sq|)/|pq|$ and conclude,

$$\frac{\delta_G(u, v)}{|uv|} = \frac{|us| + |sv| - \epsilon}{|uv|} \leq \frac{|ps| + |sq|}{|pq|} - \frac{\epsilon}{|uv|} \leq \frac{|ps| + |sq|}{|pq|} - \frac{\epsilon}{|pq|} = t.$$

If \overline{uv} is parallel to \overline{pq} , then due to the similarity of the triangles $\triangle(usv)$ and $\triangle(psq)$ we have $(|us| + |sv|)|uv| = (|ps| + |sq|)|pq|$.

Otherwise, assume w.l.o.g. that $|up| < |vq|$ and let v' be a point on $\overline{qq'}$ such that uv' is parallel to \overline{pq} . Note that $\angle(uvv') > \pi/2$ and thus, by Lemma 4 we have, $(|us| + |sv|)/|uv| \leq (|us| + |sv'|)/|uv'| \leq (|ps| + |sq|)/|pq|$.

2. One of $\{u, v\}$ is on the right or left side of the trapezoid and the other is on the horizontal component: Assume w.l.o.g. that u is on the left side and v is on the horizontal component. Recall our note regarding every triangle gadget base being followed by two edges of length $R/2$ on the left and thus $\delta_G(p', v) \leq (2/3 + t/3)|p'v|$. By the cosines law we have

$$|uv| \geq \sqrt{|up'|^2 + |p'v|^2 - 2|up'||p'v| \cos(\pi + \alpha_t)}.$$

(we do not use equality since v may not lie in the horizontal segment itself, but on a top vertex of a triangle gadget, and than $\angle(up'v) > \alpha_t$). Hence, we receive

$$\begin{aligned} \delta_G(u, v) &= \delta_G(u, p') + \delta_G(p', v) \\ &\leq |up'| + \left(\frac{2}{3} + \frac{t}{3}\right)|p'v| \\ &= \sqrt{|up'|^2 + \left(\frac{2}{3} + \frac{t}{3}\right)^2|p'v|^2 + 2|up'||p'v|\left(\frac{2}{3} + \frac{t}{3}\right)} \\ &\leq \sqrt{t^2|up'|^2 + t^2|p'v|^2 - 2|up'||p'v|t^2 \cos(\pi/2 + \alpha_t)} \leq t|uv| \end{aligned}$$

3. Both u and v are on the horizontal component: Let u' and v' be the projections of u and v on the horizontal segment. We have $\delta_G(u, v) \leq t|u'v'| \leq t|uv|$.
4. Both u and v are on the left or the right side of the trapezoid: We have $\delta_G(u, v) = |uv| < t|uv|$.

□

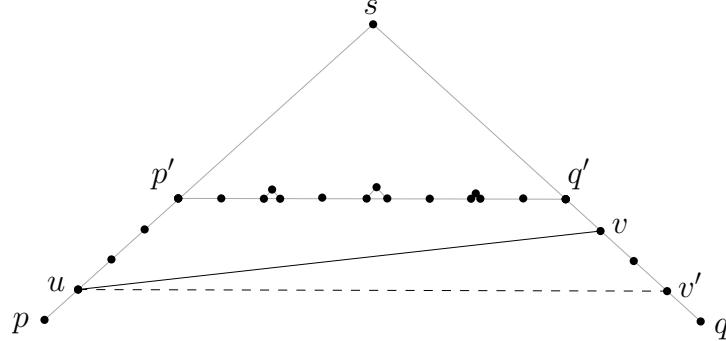


Fig. 7. Illustration of cases 1 in the proof of Lemma 5.

The only thing left to prove is that the above reduction can be computed by a Turing machine in polynomial time. Clearly, P is of polynomial size in n ; however, not all the points in P can necessarily be represented in polynomial number of bits. More precisely, setting the coordinates of the points on the horizontal segment to integer values may cause the values of other points' coordinates to have large representation size. This issue can be easily solved by rounding the values of the coordinates of the rest of the points to have up to n digits after the decimal point as follows. The y-coordinates of the triangle gadgets' top vertices are rounded down (note that the x-coordinate has polynomial size representation), and for the points on the sides, the rounding is done in such a way that increases the top angles of the trapezoid and does not increase their length.

Note that as n increases, the efficiency of the gadget t -shortcuts increases and tends to its initial value, while the efficiency of the other t -shortcuts does not exceed its initial value. Hence, for sufficiently large values of n , the absence of a proper partition of the reduction input X guarantees that a t -spanner for P of weight at most w does not exist. In case a proper partition of X exists, the suggested spanner over the output set P is still of weight smaller than w and has stretch factor at most t . Thus, for sufficiently large values of n , the reduction remains correct after the rounding. Obviously, the PARTITION problem remains NP-hard after restricting n to be larger than some constant n_0 .

In conclusion, after presenting proper reductions for all values of t , the main theorem follows.

Theorem 1. *The decision problems LWPSt and LWPSt, and the optimization problems MWSt, MWPSt, and MDG are NP-hard.*

Although all the problems addressed in Theorem 1 refer to the restricted case where the underlying graph is the complete Euclidean graph, our results, obviously, apply to the cases where the underlying graph is a general geometric graph (not necessarily the complete graph) and where the underlying graph is a general weighted graph (not necessarily geometric). This is stated in the following corollary.

Corollary 2. *The variations of the decision problems $LWSt$ and $LWPSt$ and the optimization problems $MWSt$, $MWPSt$, and MDG addressing a geometric or general weighted underlying graph are NP-hard.*

References

1. P. Bose, P. Carmi, M. Farshi, A. Maheshwari, and M. H. M. Smid. Computing the greedy spanner in near-quadratic time. In *SWAT*, pages 390–401, 2008.
2. L. Cai. Np-completeness of minimum spanner problems. *Discrete Applied Mathematics*, 48(2):187–194, 1994.
3. L. Cai and D. G. Corneil. Tree spanners. *SIAM J. Discrete Math.*, 8(3):359–387, 1995.
4. B. Chandra, G. Das, G. Narasimhan, and J. Soares. New sparseness results on graph spanners. *International Journal of Computational Geometry and Applications*, 5:125–144, 1995.
5. G. Das, P. J. Heffernan, and G. Narasimhan. Optimally sparse spanners in 3-dimensional euclidean space. In *Symposium on Computational Geometry*, pages 53–62, 1993.
6. G. Das and G. Narasimhan. A fast algorithm for constructing sparse Euclidean spanners. *Int. Journal on Computational Geometry and Applications*, 7(4):297–315, 1997.
7. G. Das, G. Narasimhan, and J. Salowe. A new way to weight malnourished Euclidean graphs. In *Proceedings of the 6th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 215–222, 1995.
8. J. Gudmundsson and M. H. M. Smid. On spanners of geometric graphs. *Int. J. Found. Comput. Sci.*, 20(1):135–149, 2009.
9. R. Klein and M. Kutz. Computing geometric minimum-dilation graphs is np-hard. In *Graph Drawing*, pages 196–207, 2006.
10. G. Narasimhan and M. Smid. *Geometric Spanner Networks*. Cambridge University Press, New York, NY, USA, 2007.